

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	6138	"I/O operation"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/01/23 21:54
L2	31	bypass\$3 with L1	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/01/23 21:59
L3	2	("6754734").PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/23 21:59
L4	896	(710/5).CCLS.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/23 22:05
L5	299	(710/6).CCLS.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/23 22:05
L6	98	(710/24).CCLS.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/23 22:05
L7	2867	(707/100).CCLS.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/23 22:05
L8	1623	(707/101).CCLS.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/23 22:05

L9	1807	(707/200).CCLS.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/01/23 22:05
----	------	-----------------	---	----	-----	------------------

**IEEE Xplore®**
RELEASE 1.8Welcome
United States Patent and Trademark Office
IEEE Xplore®
1 Million Documents
1 Million Users
...And Growing
» Search Results[Help](#) | [FAQ](#) | [Terms](#) | [IEEE Peer Review](#)[Quick Links](#)**Welcome to IEEE Xplore®**

- ☐ Home
- ☐ What Can I Access?
- ☐ Log-out

Tables of Contents

- ☐ Journals & Magazines
- ☐ Conference Proceedings
- ☐ Standards

Search

- ☐ By Author
- ☐ Basic
- ☐ Advanced
- ☐ CrossRef

Member Services

- ☐ Join IEEE
- ☐ Establish IEEE Web Account
- ☐ Access the IEEE Member Digital Library

IEEE Enterprise

- ☐ Access the IEEE Enterprise File Cabinet

Print Format[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#) | [Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#) | [No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

Copyright © 2004 IEEE — All rights reserved

Your search matched **1** of **1117580** documents.A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance** in **Descending** order.**Refine This Search:**

You may refine your search by editing the current search expression or entering a new one in the text box.

☐ Check to search within this result set**Results Key:****JNL** = Journal or Magazine **CNF** = Conference **STD** = Standard**1 Scalable resource management in high performance computers***Frachtenberg, E.; Petrini, F.; Fernandez, J.; Coll, S.;*

Cluster Computing, 2002. Proceedings. 2002 IEEE International Conference on , 23-26 Sept. 2002

Pages:305 - 314

[\[Abstract\]](#) [\[PDF Full-Text \(330 KB\)\]](#) **IEEE CNF**



US Patent & Trademark Office

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: ☒ The ACM Digital Library ☐ The Guide

I/O bypassing

SEARCH

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used **I/O bypassing**Found **21,096** of **148,786**Sort results
byrelevance ☒Display
resultsexpanded form ☒[Save results to a Binder](#)[Search Tips](#)☐ Open results in a new windowTry an [Advanced Search](#)Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐**1** [Improving I/O performance with a conditional store buffer](#)

Lambert Schaelicke, Al Davis

November 1998 **Proceedings of the 31st annual ACM/IEEE international symposium on Microarchitecture**Full text available: [pdf\(2.53 MB\)](#)Additional Information: [full citation](#), [references](#), [index terms](#)**2** [Systems and techniques: Basic I/O handling on burroughs B6500](#)

Rajani M. Patel

October 1969 **Proceedings of the second symposium on Operating systems principles**Full text available: [pdf\(848.53 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#)

The approach to processing basic Input/Output in B6500 hardware design and software implementation is discussed in this paper. Hardware I/O structure necessary to the understanding of the approach is described first. The representation of the I/O queue and the algorithms used in handling I/O requests are described in detail to emphasize the ease with which the I/O handling portions of the executive system may be modified to suit any installation. Some of the I/O tables for coordinating I/O activ ...

3 [Integrating parallel file I/O and database support for high-performance scientific data management](#)

Jaechun No, Rajeev Thakur, Alok Choudhary

November 2000 **Proceedings of the 2000 ACM/IEEE conference on Supercomputing (CDROM)**Full text available: [pdf\(174.64 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)[Publisher Site](#)

Many scientific applications have large I/O requirements, in terms of both the size of data and the number of files or data sets. Management, storage, efficient access, and analysis of these data present an extremely challenging task. Traditionally, two different solutions are used for this problem: file I/O or databases. File I/O can provide high performance but is tedious to use with large numbers of files and large and complex data sets. Databases can be convenient, flexible, and powerf ...

4 [Exploiting early sorting and early partitioning for decision support query processing](#)

J. Claussen, A. Kemper, D. Kossmann, C. Wiesner

December 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 9 Issue 3Full text available: [pdf\(478.23 KB\)](#)Additional Information: [full citation](#), [abstract](#), [index terms](#)

Decision support queries typically involve several joins, a grouping with aggregation, and/or

sorting of the result tuples. We propose two new classes of query evaluation algorithms that can be used to speed up the execution of such queries. The algorithms are based on (1) *early sorting* and (2) *early partitioning*— or a combination of both. The idea is to push the sorting and/or the partitioning to the leaves, i.e., the base relations, of the query evaluation plans (QEPs) and ...

Keywords: Decision Support Systems, Early sorting and partitioning, Hash joins and hash teams, Performance evaluation, Query processing and optimization

5 Data organization and I/O in a parallel ocean circulation model

Chris H. Q. Ding, Yun He

January 1999 **Proceedings of the 1999 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available:  [pdf\(90.16 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

6 Implementation of collective I/O in the Intel Paragon parallel file system: initial experiences

Rajesh Bordawekar

July 1997 **Proceedings of the 11th international conference on Supercomputing**

Full text available:  [pdf\(1.13 MB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#)

7 An I/O network architecture of the distributed shared-memory massively parallel computer JUMP-1

Hironori Nakajo, Satoshi Ohtani, Takashi Matsumoto, Masadi Kohata, Kei Hiraki, Yukio Kaneda

July 1997 **Proceedings of the 11th international conference on Supercomputing**

Full text available:  [pdf\(1.32 MB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#)

8 Evaluation of the Raw Microprocessor: An Exposed-Wire-Delay Architecture for ILP and Streams

March 2004 **ACM SIGARCH Computer Architecture News , Proceedings of the 31st annual international symposium on Computer architecture**, Volume 32 Issue 2

Full text available:  [pdf\(376.05 KB\)](#) Additional Information: [full citation](#), [abstract](#)

This paper evaluates the Raw microprocessor. Raw addresses the challenge of building a general-purpose architecture that performs well on a larger class of stream and embedded computing applications than existing microprocessors, while still running existing ILP-based sequential programs with reasonable performance in the face of increasing wire delays. Raw approaches this challenge by implementing plenty of on-chip resources - including logic, wires, and pins - in a tiled arrangement, and exposing the ...

9 Operating system support for high-speed communication

Peter Druschel


September 1996 **Communications of the ACM**, Volume 39 Issue 9

Full text available:  [pdf\(313.01 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

10 Disk cache—miss ratio analysis and design considerations

Alan J. Smith

August 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 3

Full text available:  [pdf\(3.13 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The current trend of computer system technology is toward CPUs with rapidly increasing processing power and toward disk drives of rapidly increasing density, but with disk performance increasing very slowly if at all. The implication of these trends is that at some point the processing power of computer systems will be limited by the throughput of the input/output (I/O) system. A solution to this problem, which is described and evaluated in this paper, is disk cache

11 Gathering at the well: creating communities for grid I/O

Douglas Thain, John Bent, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau, Miron Livny
November 2001 **Proceedings of the 2001 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available:  [pdf\(139.41 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Grid applications have demanding I/O needs. Schedulers must bring jobs and data in close proximity in order to satisfy throughput, scalability, and policy requirements. Most systems accomplish this by making either jobs or data mobile. We propose a system that allows jobs and data to meet by binding execution and storage sites together into *I/O communities* which then participate in the wide-area system. The relationships between participants in a community may be expressed by the ClassAd ...

12 TFLOPS PFS: architecture and design of a highly efficient parallel file system

Sharad Garg
November 1998 **Proceedings of the 1998 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available:  [html\(39.91 KB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

In recent years, many commercial Massively Parallel Processor (MPP) systems have been available to the computing community. These systems provide very high processing power (up to hundreds of GFLOPs), and can scale efficiently with the number of processors. However, many scientific and commercial applications that run on these multiprocessors may not experience significant benefit in terms of speedup and are bottlenecked by their I/O requirements. Although these multiprocessors may be configured ...

Keywords: I/O nodes, PFS, TFLOPS, compute nodes, fyod, parallel I/O, parallel file server

13 Functional-join processing

R. Braumandl, J. Claussen, A. Kemper, D. Kossmann
February 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 8 Issue 3-4

Full text available:  [pdf\(486.22 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

Inter-object references are one of the key concepts of object-relational and object-oriented database systems. In this work, we investigate alternative techniques to implement inter-object references and make the best use of them in query processing, i.e., in evaluating functional joins. We will give a comprehensive overview and performance evaluation of all known techniques for simple (single-valued) as well as multi-valued functional joins. Furthermore, we will describe special *order-preser* ...

Keywords: Functional join, Logical OID, Object identifier, Order-preserving join, Physical OID, Pointer join, Query processing

14 An executive system implemented as a finite-state automaton

Roy E. Heistand
November 1964 **Communications of the ACM**, Volume 7 Issue 11

Full text available:  [pdf\(1.05 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The 473L command and control system used by the Air Force permits many operators to access large data files through the use of a computer. The man-machine interface is

satisfied by several communication consoles from which operators may enter queries and view replies. A data link permits remote stations to send messages, status reports and inventories directly to the computer. The information received over the online data link is used to update the data files which are stored on disk. ...

15 Session 4: innovative solutions: Survival by defense-enabling

Partha Pal, Franklin Webber, Richard Schantz

September 2001 **Proceedings of the 2001 workshop on New security paradigms**

Full text available:  [pdf\(783.75 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Attack survival, which means the ability to provide some level of service despite an ongoing attack by tolerating its impact, is an important objective of security research. In this paper we present a new approach to survivability and intrusion tolerance. Our approach, which we call "survival by defense" is based on the observation that many applications can be given increased resistance to malicious attack even though the environment in which they run is untrustworthy. This paper describes the ...

16 The characterization of two scientific workloads using the CRAY X-MP performance monitor

Elizabeth Williams, C. Thomas Myers, Rebecca Koskela

November 1990 **Proceedings of the 1990 ACM/IEEE conference on Supercomputing**

Full text available:  [pdf\(1.19 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

The weekend production period on a CRAY X-MP was monitored for several months at each of two supercomputing sites. The hardware performance monitor available on the X-MP was used to collect the data at each site. Various metrics are computed using the measured data. These metrics include rates, such as mops and mflops, percent vectorization, average vector lengths, memory reference rates, memory contention, the operation mix, I/O rates, hardware metrics such as operations per clock period, and p ...

17 A server host system on the ARPANET

Robert T. Braden

September 1977 **Proceedings of the fifth symposium on Data communications**

Full text available:  [pdf\(820.79 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Computer networking is now an established technology, and its applications are expected to proliferate rapidly. As a result, the design of future operating systems should be influenced by the requirements of networking. The design of entirely-new operating systems for widely-used CPU's is a rather uncommon event, however; more often, it is necessary to add a network interface to an existing operating system. The consequent host software development can be a major undertaking, ...

18 Towards a theory of cache-efficient algorithms

Sandeep Sen, Siddhartha Chatterjee, Neeraj Dumir

November 2002 **Journal of the ACM (JACM)**, Volume 49 Issue 6

Full text available:  [pdf\(273.41 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present a model that enables us to analyze the running time of an algorithm on a computer with a memory hierarchy with limited associativity, in terms of various cache parameters. Our cache model, an extension of Aggarwal and Vitter's I/O model, enables us to establish useful relationships between the cache complexity and the I/O complexity of computations. As a corollary, we obtain cache-efficient algorithms in the single-level cache model for fundamental problems like sorting, FFT, and an i ...

Keywords: Hierarchical memory, I/O complexity, lower bound

19 Locking granularity revisited

Daniel R. Ries, Michael R. Stonebraker

June 1979 **ACM Transactions on Database Systems (TODS)**, Volume 4 Issue 2

Full text available:  [pdf\(1.22 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Locking granularity refers to the size and hence the number of locks used to ensure the consistency of a database during multiple concurrent updates. In an earlier simulation study we concluded that coarse granularity, such as area or file locking, is to be preferred to fine granularity such as individual page or record locking. However, alternate assumptions than those used in the original paper can change that conclusion. First, we modified the assumptions concerning the placement ...

Keywords: concurrency, database management, locking granularity, locking hierarchies, multiple updates

20 [Implementation techniques: The Multics input/output system](#)

R. J. Feiertag, E. I. Organick

June 1972 **ACM SIGOPS Operating Systems Review**, Volume 6 Issue 1/2

Full text available:  [pdf\(838.54 KB\)](#)

Additional Information: [full citation](#), [abstract](#)

An I/O system has been implemented in the Multics system that facilitates dynamic switching of I/O devices. This switching is accomplished by providing a general interface for all I/O devices that allows all equivalent operations on different devices to be expressed in the same way. Also particular devices are referenced by symbolic names and the binding of names to devices can be dynamically modified. Available I/O operations range from a set of basic I/O calls that require almost no knowledge ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)